

Intrusion Detection Methodologies

A White Paper

By Robert A. Clyde



AXENT Technologies, Inc.

1. The "business problem": Keeping the bad guys out

2. "Why Intrusion Detection?"

Taking advantage of "Free Stuff"
More computers than people

3. Early Intrusion Detection Efforts

Having "the answer" without solving the problem equals no answer at all

4. Intrusion Detection—Essential Functionality

Clearly Defined: "Intrusion Detection" is more than just a coded application

SWATting the problem of keeping current on new attacks

5. What is a "Network?"

6. Types of Intrusion Detection Tools

A. Post-event audit trail analysis

B. Real-time packet analysis

A new "Business Problem:" more point solutions without looking at the whole network

C. Real-time activity monitoring

7. Comparison of Detection Methods

8. Conclusion

1. The "business problem": *Keeping the bad guys out*

Internet and internal network attacks on corporate enterprises seem inescapable in today's computing environment. Most companies admit to having been attacked over the past year. While the most costly attacks have been from the inside, external attacks from hackers and competitors are rising dramatically. How do you know when you're under attack? Chances are you can already create enough audit trail data, but who has time to look at it?

Intrusion Detection tools solve this problem by automatically discovering and responding to attacks. This paper investigates the need for Intrusion Detection, discusses lessons learned from early Intrusion Detection efforts, and explores the different types of Intrusion Detection tools available. The paper compares and contrasts the three common methodologies used for Intrusion Detection and discusses the advantages and disadvantages inherent to various architectures.

2. "Why Intrusion Detection?"

The 1997 annual Ernst & Young security survey indicated that 46% of the respondents considered intrusions a major concern. This rose dramatically from 16% in 1996. U.S. government penetration tests at the Department of Defense over the last two years showed that less than 4% of the systems broken into were able to detect the attack. Even more disturbing, less than 1% took any response.


Taking advantage of "Free Stuff"

A few years ago, hacking took a lot of time and study. While expert hackers still abound, the Internet has entered a new era. Using almost any search engine, average Internet users can quickly find information describing how to break into systems; for example, simply searching for key words like *hacking*, *password cracking*, and *Internet security*. Thousands of sites publish step-by-step instructions as to how to break into Windows NT systems, Web Servers, UNIX systems, etc. The sites often include tools that automate the hacking process. In many cases the tools have easy to use graphical interfaces. For instance, a tool called "crack" automatically attempts to guess UNIX passwords. A similar tool called L0phtcrack breaks Windows NT passwords. A software probe called SATAN discovers vulnerable systems in a network and reports on the specific holes that can be exploited.

What does all this mean? Almost anyone with the motivation to break into systems can quickly obtain the technology to do so without having to become an expert hacker.

Attacks come from both the inside and the outside. As the survey in the following chart illustrates, disgruntled employees actually represent a larger threat and typically cause more damage than hacker attacks. An effective Intrusion Detection solution should detect attacks from both inside and outside the network.

Sources of Security Threats



More computers than people

With the explosion of Internet connectivity and the pervasive access every day users have to both internal and external networks, experts have seen a tremendous rise in attacks and corporate and government networks. At the same time the complexity of our enterprises has increased rapidly. Many organizations report that they have more computer systems than users. Add to this the diversity of operating system platforms, routers, network protocols, applications, web servers, databases, etc., and we can quickly see why trying to spot an attack becomes extremely difficult. Without sophisticated tools, it's nearly impossible.

Nevertheless, nearly every organization wants to know when they are under attack. Enter Intrusion Detection technology. Intrusion Detection tools automatically detect attacks and threats and ideally provide some type of response.

3. Early Intrusion Detection Efforts

In the early 1980s, conventional wisdom dictated that the best way to detect intrusions was to create logs or audit trails of all security relevant activity. As a result most operating systems, databases, routers, and mission-critical applications generate audit trails. The original idea was that a security administrator would review the audit logs looking for suspicious events. This seemed like a fine idea when companies only had a few systems and a few users.

The industry quickly realized that no one had time to read all that audit trail data. So a few enterprising developers built query and reporting programs to help analyze the audit trail in an attempt to find trouble spots. For example, in 1984, Clyde Digital Systems developed a product called AUDIT, which automatically searches through OpenVMS audit trails looking for suspicious events (incidentally, that product is still in use today). In 1987, a U.S. Government-funded project called IDES at Stanford Research Institute read audit trails and created profiles of normal use patterns for users and then reported deviations.

Having "the answer" without solving the problem equals no answer at all

Intrusion Detection efforts throughout the 1980's and early 90's tended to focus on post-event audit trail analysis. Most companies, however, did not make use of such tools. Unfortunately, as the number of users, systems, applications, and databases grew, so did the audit trails now grow so large that they actually can cause denial of service problems from using up too much disk space. Many production environments routinely turn off audit trails to avoid disruptions to production systems.

So the current situation at most sites is that they plan to rely on audit trails to detect intrusions, but without staffing to review the audit trails, these sites turn off the audit trails to improve productivity. No wonder most attacks go undetected. Nobody's looking.

4. Intrusion Detection—Essential Functionality

The term "Intrusion Detection" implies discovering attacks and threats throughout an enterprise, and responding to those discoveries. Some of the automated responses typically include notifying a security administrator via a console, e-mail, pager; stopping the offending session; shutting the system down; turning off down Internet links; disabling users; or executing a predefined command procedure.

Clearly Defined: *"Intrusion Detection" is more than just a coded application*

An effective Intrusion Detection system needs to limit false positives—incorrectly identifying an attack when there is none. At the same time it needs to be effective at catching attacks. Figuratively speaking, Intrusion Detection is like a surveillance camera and alarm system all rolled into one. False alarms are distracting and reduce the effectiveness of an Intrusion Detection system. Failing to catch a break-in reduces its value even further. To detect new types of attacks an Intrusion Detection tool must have a way to be quickly updated. This is particularly challenging since updates of attack detection scenarios need to be more frequent than typical product release upgrade cycles of three to nine months. In fact, to be effective probably requires updating the software to new detection procedures on a regular basis.

SWATting the problem of keeping current on new attacks

AXENT's Information Security SWAT Team illustrates one way to address this challenge of rapid deployment of new attack scenarios. The SWAT team researches new attack techniques and security threats and tests them in the lab. It develops new Intrusion Detection scenarios in response and publishes both a description of the attack and the scenarios on an Internet web site, www.axent.com/swat/swat.htm. Customers can download and quickly deploy new Intrusion Detection scenarios every week or two.

5. What is a "Network?"

Although this may seem strange, but let's clearly define the term "network." Why? Many intrusion detection products on the market claim to be network-based, when in fact, they are only link-based packet-sniffers and analyzers. Remembering basic geometry, a network is an assembly of "nodes" and "links." You might have seen the following illustration used to define the term "network."

In the example, to meet our basic definition of a network, the illustration required single points, connected by individual lines. The points, we described as "nodes" and the lines connecting between these nodes we referred to as "links." (Individual links can connect multiple nodes as shown by the middle link in the picture, which connects three nodes. Ethernet is an example of a network link that can connect multiple nodes to a single segment.)

In the Intrusion Detection industry, much attention has been focused on the individual links, or on the individual nodes (some times referred to as "hosts"). The following section examines the various methods that the leading vendors consider as their solution to "Network-wide Security."

6. Types of Intrusion Detection Tools

As recently as the last couple of years a number of Intrusion Detection products have appeared on the market. The Intrusion Detection market is relatively new, but growing fast. Based on their underlying methodologies, today's Intrusion Detection products fall into three basic categories:

- . Post-event audit trail analysis
- B. Real-time packet analysis
- C. Real-time activity monitoring

Each of these categories has value and particular advantages and disadvantages which this paper explores. Although no single product currently falls into more than two categories, we expect the lines between Intrusion Detection products to blur. In the future, superior products will cover all three categories.

A. Post-event audit trail analysis

As mentioned earlier traditional Intrusion Detection has been to perform post-event audit trail analysis. SAIC's CMDS product and TIS' Stalker product fall into this category since they analyze certain UNIX audit trails for suspicious activity. Products in this category have enhanced value if they also perform automated audit trail and management. This means extracting and archiving critical information and purging out old and unneeded audit trail data. For instance, in addition to being a real-time activity monitor, AXENT's Intruder Alert product also performs audit trail analysis, reduction, and management.

This type of product has two key advantages. One is that it addresses the tremendous difficulties that organizations experience examining and managing audit trails. Many times the purchase of such a product can be justified on the cost savings achieved through the centralization and automation of audit trail management.

The second advantage is that investigators can go back in time and do historical analysis of events that have occurred in the past. More sophisticated products can graph results and show trend analysis by attack category, system, type of system, etc. This is particularly useful in investigations of break-ins that have taken place over a period of time.


From a network-based security perspective, the disadvantage with a pure "after-the-fact" product is that by the time it detects the security problem, it's generally too late to respond and protect the data, and the resulting consequences of the attack go far deeper into the network without resistance. Ultimately, the damage is already done by the time you find out. Also, since most hackers learn how to cover their tracks by tampering with audit trails, after-the-fact analysis often misses attacks.

B. Real-time packet analysis

Before the last couple of years, the idea of doing real-time network-wide Intrusion Detection seemed unfeasible. Systems lacked the necessary speed to perform this type of analysis without causing unacceptable performance problems. Now the market has a number of products that detect attacks in real-time and respond immediately, hopefully before damage is done.

One method of real-time Intrusion Detection is to dedicate a system to sniffing packets traveling across a single network segment (remember the "link" discussed in the Network example). And usually, the only links that could be monitored were ethernets. Using this methodology, the Intrusion Detection software is placed on the system, which puts the ethernet card in "promiscuous mode" so that the software can read and analyze all traffic. It does this by examining both the packet header fields and packet contents. The Intrusion Detection software includes an engine, which looks for specific types of network attacks, such as IP spoofing and packet floods. When the packet analysis software detects a potential problem it reponds immediately by notifying a console, beeping a pager, sending an e-mail, or even shutting down the network session. This category includes products such as Wheelgroup's NetRanger, ISS' RealSecure, and Network Associates' CyberCop.

The diagram below shows a typical deployment of sniffers for doing packet analysis. A sniffer is placed outside the firewall to detect attack attempts coming from the Internet. A sniffer is also placed inside the network to detect Internet attacks, which penetrate the firewall and to assist in detecting internal attacks and threats. For full enterprise coverage sniffers must be placed on each network segment. Also, more sophisticated tools should be able to remotely manage the various sniffers, collate the information gathered, and display the enterprise-wide information on a console.



The advantages of the packet analysis technique are that there are certain network-oriented attacks (IP spoofing, packet storms, etc.) that are best detected via packet examination. In addition, you do not need to put software on various hosts throughout the network. But remember the basic definition of a network: an organization of nodes and links. A packet analyzer monitors traffic on the links, but does not monitor the nodes, which are key pieces of any network. Calling a packet analyzer "Network-based" Intrusion Detection ignores the basic definition of a network—which includes nodes as well as links.

Using the packet-sniffing methodology as the exclusive Intrusion Detection technique has other disadvantages as well.

1. Packet analysis Intrusion Detection is distant from the mission-critical applications and data it is trying to protect. It's a little bit like trying to perform surveillance on a bank vault by sitting on a hill a half mile away and watching who goes in and out of the bank with binoculars. You might detect some suspicious characters and take action, but you really don't know what's actually happening inside the bank vault.
2. Packet analysis does not detect typical attacks like:
 - Exploiting a buffer overflow flaw on UNIX to gain root
 - Browsing for files that the user shouldn't have access to
 - Attacking mission-critical servers through dial-up lines
 - Improperly modifying firewall or router settings
 - Inserting Trojan horses on systems, such as changing the Windows NT log-in program
 - Illegally using a mission critical application (e.g., funds transfer system)
 - Tampering with the content of web pages and a web server
 - Exploiting a Windows NT registry vulnerability to gain administrator access
 - Inappropriately accessing a database
1. Sniffers require dedicated hardware for network each segment being monitored. The cost of the hardware increases depending upon the speed of the network link. The sniffer box must also be capable of keeping up with the volume of traffic. As faster networks are deployed this will require significant hardware upgrades for the packet analyzers.
2. Packet sniffers do little in the space of encrypted packets. At best, sniffers can acknowledge that a packet transferred across the link. But since the data is encrypted, the sniffer cannot report in context as to what the packet contained.

A new "Business Problem:" more point solutions without looking at the whole network

Companies have already spent vast amounts of money on products that analyze and filter packets, like routers and firewalls. By analyzing the audit trails on routers and firewalls it is possible to detect most of the

network-oriented attacks that a packet analysis Intrusion Detection product would find. How many separate boxes for security on a given network segment are companies going to buy? First they buy a router, then a firewall, and now an Intrusion Detection packet analyzer. All of these devices analyze packets. Many organizations will balk at the extra expense for seemingly redundant hardware. Since the firewall looks at all of the packets any way, shouldn't it detect attacks? Nevertheless, some security experts argue that the packet analyzer does provide an additional layer of separate protection beyond the firewall, *but only on a limited basis*.

C. Real-time activity monitoring

An effective method for real-time Intrusion Detection is to monitor security-related activity occurring on the various systems and devices that make up the network. To refer back to the bank analogy, this is like putting surveillance cameras and alarms inside the bank watching the tellers and vault. The surveillance is close to the valuables. While most activity monitors watch the operating system audit trails, more sophisticated tools also

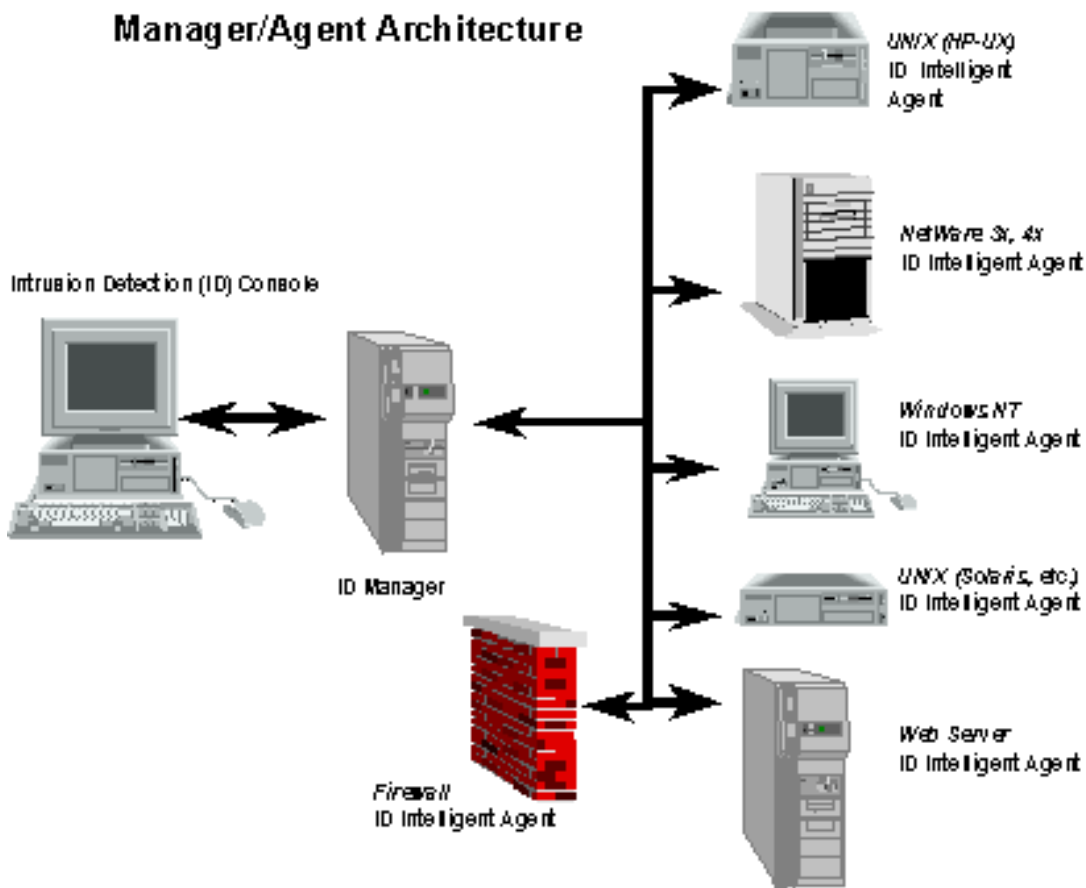
- Track audit trails from applications, databases, web servers, routers, firewalls, etc.
- Monitor critical files for Trojan horses, unauthorized changes, etc.
- Watch TCP and UDP port activity
- Accept SNMP traps and triggers

Real-time activity monitors can detect attacks such as attempts to access unauthorized sensitive files or to replace the log-in program with a new version. Unlike packet sniffers they can detect when a user illegally obtains "root" or administrator access. When suspicious activity is detected the real-time activity monitor can take immediate action before damage is done. This action typically includes notifying a console, sending e-mail, beeping a pager, disabling a user account, terminating the intruder's process, terminating the intruder's session, shutting the system down, or executing a command procedure.

The advantage of real-time activity monitors is that they deploy close to the mission-critical data and applications. Monitoring for attacks from both the inside and the outside the network becomes much easier, since all of the devices are being watched. In addition, many application and operating system-level attacks are not discernable at the packet level and require system level monitoring to detect.

Activity monitors include two basic categories: single system and manager/agent. A single system activity monitor runs on only one system in the network and detects intrusions based solely on what it finds on that system. Such an Intrusion Detection tool is installed and executed on a stand-alone, system-by-system basis. For Windows NT, Intrusion Detection, Inc.'s Kane Security Monitor and TIS' WebStalker are examples of stand-alone, single-system tools.

A manager/agent Intrusion Detection solution has agents covering systems and network devices throughout the enterprise. These agents are connected to various managers, which are in turn connected to an Intrusion Detection console. From the console you can remotely install and upgrade agents, define Intrusion Detection scenarios across agents, and track intrusions as they occur throughout the enterprise. The diagram below shows the architecture of a manager/agent solution.



An example, of a manager/agent real-time Intrusion Detection architecture is AXENT's OmniGuard/Intruder Alert. Intruder Alert runs across Windows NT, UNIX, and NetWare (more than 50 operating system versions). It also monitors audit trails from Cisco routers, web servers, and various firewalls.

Intruder Alert's manager/agent architecture offers the following advantages:

- Manages Intrusion Detection from a central console, while still monitoring activity throughout the entire network.
- Relies on the devices themselves for first-level packet monitoring. Events that manage to slip through the device's capabilities to catch them are then evaluated by Intruder Alert.
- Correlates suspicious activity as it occurs in multiple locations in the network. For example, an intruder may use a hacker program to attempt to guess the root password on a hundred UNIX systems at the same time.
- Quickly updates the various agents in the network with new attack scenarios. The vendor could publish these scenarios on the web so that customers could then download them and rapidly deploy them throughout the enterprise.
- Detects intrusions even if network connections are encrypted or if attackers use direct dial-up connections.
- Logs critical security activity on manager systems. This makes it difficult for hackers to cover their tracks since activity is logged on another system in the network, not just a local audit trail. It also centralizes and facilitates audit trail management.

7. Comparison of Detection Methods

The chart below shows a brief comparison of the basic features of the various methods of Intrusion Detection. The final section of the chart shows what types of security threats and attacks each method can detect. A check mark means that it can detect and respond. A "d" means it can only detect, but that it can't provide an immediate

response. The presence of a "d" or checkmark does not necessarily mean that a particular product currently has the indicated function, but that the methodology makes supporting that function straightforward.

Intrusion Detection Functionality	After-the-fact audit trail analysis	Real-time packet analysis	Real-time activity monitoring	
			Stand-alone	Manager/ Agent
Features				
Historical audit trail analysis	d		Ö	Ö
No dedicated hardware required	d		Ö	Ö
Real-time attack detection		Ö	Ö	Ö
Immediate response		Ö	Ö	Ö
Integration with system & network management framework	d	Ö	Ö	Ö
Network-wide reporting	d	Ö		Ö
Collating activity across network				Ö
Handle O/S specific events (e.g., Windows NT, UNIX, NetWare)	d		Ö	Ö
Watch application, database, web server, etc. security activity	d		Ö	Ö
Automate distribution of ID software through network				Ö
Automatically update attack scenarios from central console	d	Ö		Ö
Security Attacks and Threats Detected				
Password cracking attempts	d		Ö	Ö
Password guessing on multiple systems (one password at a time)				Ö
IP Spoofing		Ö		Ö
SYN Flood		Ö	Ö	Ö
Land Attack		Ö		
Attacks through dial-up modems	d		Ö	Ö

Attack from inappropriate IP address		Ö		Ö
Illegal "Root" grabbing			Ö	Ö
Critical file tampering	d		Ö	Ö
Trojan horse detection	d		Ö	Ö
Browsing files (snooping)	d		Ö	Ö
Snooping across multiple systems				Ö
Response Types				
Alert central console		Ö		Ö
Send e-mail		Ö	Ö	Ö
Send message to pager		Ö	Ö	Ö
Disable intruder's user account			Ö	Ö
Terminate network access		Ö		Ö
Terminate intruder's session		Ö	Ö	Ö
Shutdown system			Ö	Ö
Terminate intruder's user process			Ö	Ö
Generate SNMP Trap		Ö	Ö	Ö
Record event on security server		Ö		Ö
Execute command procedure		Ö	Ö	Ö

The previous chart clearly shows that while all Intrusion Detection methodologies are useful, manager/agent real-time activity monitoring has the most flexible architecture. It can pick up information from routers, firewalls, and other sources to detect many different kinds of attacks.

8. Conclusion

Intrusion detection is critical in today's complex enterprises. Attempting to manually review audit trails is hopelessly time-consuming and a losing battle given the number of systems and different types of audit trails. Today's enterprises need automated Intrusion Detection tools. These tools fall into three categories, post-event audit trail analysis, real-time packet analysis, and real-time activity monitoring.

All three types of Intrusion Detection methods have merit, although post-event monitoring lacks the capability for immediate response to avoid or reduce damage. Real-time packet analysis is interesting for detecting certain low-level packet attacks, but is too far from the system—and does not effectively solve the network-wide Intrusion Detection problem alone. Real-time activity monitoring that considers both host and link activity seems the appropriate solution for Intrusion Detection.

A manager/agent architecture provides the ability to monitor intrusions network-wide and to perform audit trail analysis and management as well as real-time Intrusion Detection. This covers both the first and third methods. Hooking packet analysis into a manager/agent architecture is really just a special case of adding a new type of agent to the manager/agent product.

As Intrusion Detection moves into the future we expect to see specific products that span all three types of types of Intrusion Detection. Because of the enabling infrastructure they already possess, products with a manager/agent architecture, like Intruder Alert, are most likely to adequately focus on all three of the Intrusion Detection methodologies.